



(11) **EP 2 958 077 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:  
**23.12.2015 Bulletin 2015/52**

(51) Int Cl.:  
**G06T 7/00 (2006.01)**

(21) Application number: **15172032.3**

(22) Date of filing: **15.06.2015**

(84) Designated Contracting States:  
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB  
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO  
PL PT RO RS SE SI SK SM TR**  
Designated Extension States:  
**BA ME**  
Designated Validation States:  
**MA**

(72) Inventors:  
• **Reso, Matthias**  
**30177 Hannover (DE)**  
• **Jachalsky, Jörn**  
**30974 Wennigsen (DE)**  
• **Rosenhahn, Bodo**  
**30952 Ronnenberg (DE)**


(30) Priority: **20.06.2014 EP 14305959**  
**28.07.2014 EP 14306209**

(74) Representative: **Schmidt-Uhlig, Thomas**  
**Deutsche Thomson OHG**  
**European Patent Operations**  
**Karl-Wiechert-Allee 74**  
**30625 Hannover (DE)**

(71) Applicant: **Thomson Licensing**  
**92130 Issy-les-Moulineaux (FR)**

(54) **METHOD AND APPARATUS FOR GENERATING TEMPORALLY CONSISTENT SUPERPIXELS**

(57) A method and an apparatus for generating superpixels for a sequence of images are described. A cluster assignment generator (22) generates (10) a cluster assignment for a first image of the sequence of images, e.g. by clustering pixels of the first image into superpixels or by retrieving an initial cluster assignment for the first image and processing only contour pixels with regard to their cluster assignment. A label propagator (23) initializes (11) subsequent images based on a label propagation using backward optical flow. A contour pixel processor (24) then processes (12) only contour pixels with regard to their cluster assignment for subsequent images of the sequence of images.



**Fig. 5**

**Description****FIELD OF THE INVENTION**

5 **[0001]** The present invention relates to a method and an apparatus for generating temporally consistent superpixels for a video sequence. More specifically, a method and an apparatus for generating temporally consistent superpixels are described, which make use of contour evolution.

**BACKGROUND OF THE INVENTION**

10 **[0002]** Superpixel algorithms represent a very useful and increasingly popular preprocessing step for a wide range of computer vision applications, such as video segmentation, tracking, multi-view object segmentation, scene flow, 3D layout estimation of indoor scenes, interactive scene modeling, image parsing, and semantic segmentation. Grouping similar pixels into so called superpixels leads to a major reduction of the image primitives. This results in an increased  
15 computational efficiency for subsequent processing steps, allows for more complex algorithms computationally infeasible on pixel level, and creates a spatial support for region-based features.

**[0003]** Superpixel algorithms group pixels into superpixels. As indicated in [1], superpixels are local, coherent, and preserve most of the structure necessary for segmentation at the scale of interest. Superpixels should be roughly homogeneous in size and shape. Though many superpixel approaches mostly target still images and thus provide only  
20 a limited or no temporal consistency at all when applied on video sequences, some approaches target video sequences [2][3]. These approaches start to deal with the issue of temporal consistency.

**[0004]** The superpixel generation in itself does not necessarily lead to spatially coherent superpixels. Thus, a post-processing step is required to ensure the spatial connectivity of the pixels comprised in the clusters and thus the superpixels. In addition, in [4] it was stated that the post-processing method proposed in [5] assigns the isolated superpixel  
25 fragments to arbitrary neighboring superpixels without considering any similarity measure between the fragments and the superpixels they are assigned to. Contour evolution approaches as proposed in [4] can overcome this drawback, often at the cost of a high number of iterations. In addition, they often focus on still images and thus leave the temporal consistency issue unsolved.

**SUMMARY OF THE INVENTION**

30 **[0005]** It is an object of the present invention to propose an improved solution for generating temporally consistent superpixels associated to images of a sequence of images.

**[0006]** According to the invention, a method for generating temporally consistent superpixels for a sequence of images  
35 comprises:

- generating a cluster assignment for a first image of the sequence of images;
  - initializing subsequent images based on a label propagation using backward optical flow; and
  - for subsequent images of the sequence of images, processing only contour pixels with regard to their cluster assignment.
- 40

**[0007]** Accordingly, an apparatus configured to generate temporally consistent superpixels for a sequence of images comprises:

- a cluster assignment generator configured to generate a cluster assignment for a first image of the sequence of images;
- a label propagator configured to initialize subsequent images based on a label propagation using backward optical flow; and
- a contour pixel processor configured to process only contour pixels with regard to their cluster assignment for  
50 subsequent images of the sequence of images.

**[0008]** Also, a computer readable storage medium has stored therein instructions enabling generating temporally consistent superpixels for a sequence of images, wherein the instructions, when executed by a computer, cause the computer to:

- generate a cluster assignment for a first image of the sequence of images;
  - initialize subsequent images based on a label propagation using backward optical flow; and
  - for subsequent images of the sequence of images, process only contour pixels with regard to their cluster assignment.
- 55

**[0009]** Furthermore, a computer program comprises instructions enabling generating temporally consistent superpixels for a sequence of images, wherein the instructions, when executed by a computer, cause the computer to:

- generate a cluster assignment for a first image of the sequence of images;
- initialize subsequent images based on a label propagation using backward optical flow; and
- for subsequent images of the sequence of images, process only contour pixels with regard to their cluster assignment.

**[0010]** The proposed solution introduces a contour evolution-based strategy for the clustering-based superpixel approach described in [6]. Instead of processing all pixels in the video volume during the clustering only the contour pixels are processed. Therefore, in each iteration only the contour pixels can be changed, i.e. assigned to a different cluster. The other pixels keep their previous assignment. New images entering the video volume are initialized by propagating the contours of the latest image utilizing backward flow information.

**[0011]** In one embodiment, generating the cluster assignment for the first image of the sequence of images comprises clustering pixels of the first image into superpixels.

**[0012]** In another embodiment, generating the cluster assignment for the first image of the sequence of images comprises retrieving an initial cluster assignment for the first image and processing only contour pixels with regard to their cluster assignment. The initial cluster assignment for the first image is created by the tessellation of the image using geometric shapes, which can be e.g. squares, rectangles or hexagons. All pixels covered by the same shape are assigned to the same cluster. This finally results into the initial cluster assignment.

**[0013]** A fully contour-based approach for superpixels on video sequences is proposed, which is expressed in an expectation-maximization (EM) framework, and generates superpixels that are spatially coherent and temporally consistent. An efficient label propagation using backward optical flow encourages the preservation of superpixel shapes when appropriate.

**[0014]** With the proposed approach in general no post-processing step is required to ensure spatial coherency after the clustering. At the same time the generated superpixels show a high boundary/contour accuracy and a high temporal consistency. In addition, the approach works reasonably fast due to selective processing. The generated superpixels are beneficial for a wide range of computer vision applications, ranging from segmentation, image parsing to classification etc.

**[0015]** For a better understanding the invention shall now be explained in more detail in the following description with reference to the figures. It is understood that the invention is not limited to this exemplary embodiment and that specified features can also expediently be combined and/or modified without departing from the scope of the present invention as defined in the appended claims.

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0016]**

Fig. 1 shows an example of superpixels with temporal consistency;

Fig. 2 shows a 5×5 pixel-neighborhood between two superpixels;

Fig. 3 illustrates possible variations of superpixel label propagation to new frames;

Fig. 4 explains problems occurring when propagating whole superpixels by mean optical flow;

Fig. 5 schematically illustrates a method according to the invention; and

Fig. 6 schematically illustrates an apparatus configured to perform a method according to the invention.

## DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

**[0017]** Fig. 1 shows an example of superpixels with temporal consistency. While Fig. 1(a) depicts the original image, in Fig. 1(b) a full segmentation of the video was performed and a subset of superpixels was manually selected in one image and provided with grey values for visualization. Fig. 1(c) shows the same subset after several images. The same grey value means temporal connectedness.

**[0018]** For a better understanding of the main idea, in the following the contour evolution-based approach is first briefly described for still images and then extended to video using a slightly different explanatory approach.

Short brief description for still images

**[0019]** For clustering-based approaches, pixels of an image are seen as data points in a multi-dimensional feature space, in which each dimension corresponds to a color channel or image coordinate of the pixels. Superpixels are represented by clusters in this multi-dimensional feature space and each data point can only be assigned to one cluster. This assignment finally determines the over-segmentation and thus the superpixel generation.

**[0020]** In order to find an optimal solution for this assignment problem using a fully contour evolution-based approach, an energy function  $E_{total}$  is defined, which sums up the energy  $E(n,k)$  that is needed to assign a contour pixel  $n \in \mathcal{N}_c$  to a cluster  $k \in \mathcal{K}_n$ :

$$E_{total} = \sum_{n \in \mathcal{N}_c} E(n, k),$$

where  $\mathcal{N}_c$  is the set of contour pixels in the image and  $\mathcal{K}_n$  is the set of clusters that are adjacent to the contour pixel  $n$  and to the cluster to which the contour pixel  $n$  is assigned. The energy  $E(n,k)$  can be further refined as the weighted sum of a color-difference related energy  $E_c(n,k)$  and a spatial-distance-related energy  $E_s(n,k)$ :

$$E(n, k) = (1 - \alpha)E_c(n, k) + \alpha E_s(n, k).$$

**[0021]** The energy  $E_c(n,k)$  is directly proportional to the Euclidean distances between the contour pixel  $n$  and the color center of cluster  $k$  in the chosen color space. Likewise,  $E_s(n,k)$  is proportional to the Euclidean distance of the spatial position of  $n$  and the spatial position of the center of cluster  $k$ . In order to make the results independent from the image size, the spatial distance is scaled with the factor  $1/\sqrt{|\mathcal{N}|/|\mathcal{K}|}$ , where  $|\cdot|$  is the number of elements in a set,  $\mathcal{N}$  the set of all pixels in the image and  $\mathcal{K}$  the set of all clusters in the image. With the parameter  $\alpha$  the user can steer the segmentation results to be more compact or more sensitive to fine-grained image structures. For a given number of clusters  $|\mathcal{K}|$  and a user-defined  $\alpha$ , an optimal over-segmentation in terms of energy can be determined by finding a constellation of clusters that minimizes  $E_{total}$ .

**[0022]** The assignment problem is solved by applying an iterative Lloyd's algorithm, which converges to a locally optimal solution. The clusters are initialized as evenly distributed, non-overlapping rectangles or squares, for example, in the spatial subspace. To minimize the energy term  $E_{total}$  the algorithm iterates two alternating steps, the *assignment-step* and the *update-step*. In the *assignment-step*, each contour pixel  $n$  is assigned to the cluster  $k \in \mathcal{K}_n$ , for which the energy term  $E(n,k)$  has its minimum given the set  $\mathcal{K}_n$ . The other pixels that are not contour pixels keep their assignments. Based on the assignments, the parameters of the cluster centers are re-estimated in the *update-step* by calculating the mean color and mean position of their assigned pixels. The iteration stops when no changes in the *assignment-step* are detected or a maximum number of iterations have been performed.

**[0023]** A contour pixel  $n$  is a pixel with at least one adjacent pixel that is assigned to a different cluster. The set of all contour pixels  $\mathcal{N}_c$  is a (small) subset of the set  $\mathcal{N}$ . The cluster the contour pixel  $n$  is assigned to and all clusters of the adjacent pixels, which are assigned to a different cluster than  $n$ , form the set  $\mathcal{K}_n$ .

**[0024]** In order to minimize the bias caused by a constant scan-order used to traverse the image in the assignment step, it should be changed with each iteration. Possible scan-orders are, for example, left-to-right and top-to-bottom, right-to-left and top-to-bottom, left-to-right and bottom-to-top, right-to-left and bottom-to-top.

**[0025]** As  $\mathcal{N}_c$  and  $\mathcal{K}_n$  (and not  $\mathcal{N}$  and  $\mathcal{K}$ ) are taken into consideration to determine  $E_{total}$ , the processing load is significantly reduced.

Approach for videos

**[0026]** As indicated in the beginning, superpixel algorithms group spatially coherent pixels that share the same low-

level features as e.g. color or texture into small segments of approximately same size and shape. Therefore, one important and inherent constraint is that the boundaries of the superpixels should comply with the main object boundaries present in the image. This capturing of the main object boundaries is rather important for image or video segmentation built upon the initial superpixel segmentation. Moreover, for a superpixel segmentation on video content it is necessary to capture the temporal connections between superpixels in successive images in order to achieve a consistent labeling that can be leveraged for applications like tracking or video segmentation.

**[0027]** As mentioned before, the approach described here extends the approach described in [6], entitled Temporally Consistent Superpixels (TCS). For a better understanding the main ideas of TCS shall be shortly summarized in the following,

**[0028]** In general, TCS performs an energy-minimizing clustering on a video volume utilizing a hybrid clustering strategy for a multi-dimensional feature space, which is separated into a global color subspace and multiple local spatial sub-spaces.

**[0029]** More specifically, the energy-minimizing framework used in TCS clusters pixels based on their five dimensional feature vector  $[labxy]$ , which contains the three color values  $[lab]$  in CIE-Lab-color space and the pixels coordinates  $[xy]$ . In order to capture the temporal connections between superpixels in successive images, the clustering is performed over an observation window spanning  $K$  images. The separated feature space is realized in the following way. Each cluster center representing one temporal superpixel consists of one color center for the complete observation window and multiple spatial centers with one for each image in the observation window. The underlying assumption is that a temporal superpixel should share the same color in successive images but not necessarily the same position.

**[0030]** For the observation window that is shifted along the timeline an optimal set of cluster centers  $\Theta_{opt}$  as well as a mapping  $\sigma_{opt}$  of the pixels inside the observation window to these cluster centers are obtained. Therefore, an energy function is defined, which sums up the energies necessary to assign a pixel at position  $x, y$  in image  $k$  to a cluster center  $\theta \in \Theta_{opt}$ . This assignment or mapping is here denoted by  $\sigma_{x,y,k}$ .

$$E_{total} = \sum_k \sum_{x,y} (1 - \alpha) E_c(x, y, k, \sigma_{x,y,k}) + \alpha E_s(x, y, k, \sigma_{x,y,k}).$$

**[0031]** The energy needed for an assignment is the weighted sum of a color dependent energy  $E_c(x, y, k, \sigma_{x,y,k})$  and a spatial energy  $E_s(x, y, k, \sigma_{x,y,k})$ . Both energy terms are proportional to the Euclidean distance in color space and image plane, respectively. The trade-off between color-sensitivity and spatial compactness is controlled by a weighting factor  $\alpha$ , which has a range between 0 (fully color-sensitive) and 1 (fully compact). The energy function is minimized using an iterative optimization scheme, which can be viewed as an expectation-maximization (EM) approach.

**[0032]** In the E-step (Expectation-step) of iteration  $l+1$  a new estimation of the optimal mapping, here denoted as  $\hat{\sigma}_{x,y,k}^{l+1}$ , is determined, which minimizes  $E_{total}$  based on the estimation of the optimal set of cluster center  $\hat{\Theta}_{opt}^l$  calculated in the M-step (Maximization-step) of iteration  $l$ .

**[0033]** After that, the estimation of the optimal cluster set  $\hat{\Theta}_{opt}^{l+1}$  is updated in the M-step of iteration  $l+1$  given the updated mapping by calculating the mean color and mean spatial values of the assigned pixels. The alternation of the two steps continues until the energy  $E_{total}$  drops below a specific bound or a fixed number of iterations is performed. In the hybrid clustering proposed for TCS, only the  $K_F < K$  most future images in the observation window are reassigned during the optimization. For the remaining  $K - K_F$  images the determined mapping is kept in order to preserve the color clustering found.

**[0034]** The observation window is shifted along the video volume and new images entering the observation window are initialized as described in more detail below. To ensure the spatial coherency of the generated temporally consistent superpixels a contour-based post processing step is applied.

**[0035]** Revisiting the ideas of TCS, the following two observations were made: (a) In order to achieve a higher run-time performance, the initial energy-minimizing cluster and the contour-based post processing are separated steps. (b) New images added to the observation window are initialized by projecting only the spatial centers of the preceding image into the new image. As a consequence the shape information obtained in the images before is at least partially discarded.

**[0036]** In contrast to TCS, the energy function  $E_{total}$  for the proposed approach is only optimized for pixels at a contour (or boundary) of a superpixel, so called contour pixels. A contour pixel at position  $x, y$  has at least one pixel in its 4-connected neighborhood  $\mathcal{N}_{x,y}^4$ , which is assigned to a different cluster, i.e. a temporal superpixel, or is unassigned. The occurrence of unassigned pixels and their handling is described in detail below. Moreover, the assignment of a

contour pixel can only be changed to one of the clusters of the pixels in  $\mathcal{N}_{x,y}^4$ . The E-step of the optimization is expressed as

$$\hat{\sigma}_{x,y,k}^{l+1} = \underset{\hat{\sigma}_{\tilde{x},\tilde{y},k}^l: \tilde{x},\tilde{y} \in \mathcal{N}_{x,y}^4}{\operatorname{argmin}} (1 - \alpha)E_c(x, y, k, \hat{\sigma}_{\tilde{x},\tilde{y},k}^l) + \alpha E_s(x, y, k, \hat{\sigma}_{\tilde{x},\tilde{y},k}^l) \quad \forall x, y \in C_k^l.$$

where  $C_k^l$  is the set of contour pixels after iteration step  $l$  in frame  $k$ . The optimization is done for the  $K_F$  most future images in the observation window. The M-step remains unmodified. The optimization can be terminated if there are no further assignment changes for the contour pixels or if a maximum number of iterations has been reached.

**[0037]** In addition to the description above, there are two constraints. First, an assignment change is only done if the spatial coherency of the superpixels is guaranteed. This constraint prevents that fragments of a temporal superpixel are split-off during the optimization, as shown in Fig. 2. The three subfigures show a  $5 \times 5$  pixel-neighborhood between two superpixels (light grey and dark grey). If the centered pixel (medium grey in (b)) changes its assignment, the two pixels on its right lose connection to the light grey superpixel and thus they would be split-off from the main mass (as shown exemplarily in (c)). Therefore, no assignment change is performed in situations like these.

**[0038]** Second, if a contour pixel is unassigned, it is assigned to the cluster of one of its adjacent pixels based on

$\hat{\sigma}_{x,y,k}^{l+1}$ . As a consequence, the additional post-processing step required in TCS to ensure the spatial coherency is not needed anymore and can be omitted. Although especially the first constraint limits to some extent the flexibility of how the temporal superpixels can adapt to the video content, experiments have proven that it does not have a negative impact on the segmentation accuracy.

**[0039]** As mentioned above, TCS initializes new images only by projecting the spatial centers into a new image entering the observation window. Therefore, the weighted average of the dense optical flow determined over all pixels assigned to a cluster is used for the projection. Especially for image areas with a homogeneous color or structure, this can lead to an unsteady positioning of the superpixels from image to image. Depending on the application relying on the superpixel segmentation this could be an unwanted effect. In order to overcome this problem, one solution is to transfer the shape of the superpixels to the new image to be initialized. This helps to preserve the shape information as well as the superpixel constellation obtained in previous images, as it is at least the starting for the next optimization in the EM framework.

**[0040]** There are several ways to realize such an initialization of the new images. One is the shift of the complete superpixel label using the mean optical flow, as depicted in Fig. 3(a). An alternative is the usage of a dense optical flow predicted for each pixel of the superpixel. Thus, the superpixel label is projected into the new image as shown in Fig. 3(b). These two options have the following drawback: If two superpixels projected into the new image overlap, it is at least necessary to detect this collision. In addition, it is possible that there are unassigned parts in the image that have to be initialized, e.g. if adjacent superpixels are projected away from each other. Both cases are illustrated in Fig. 4, which shows a ripped open superpixel label (light grey stripes) and overlapping areas (black), and apply in the same manner to the shifting of pixels by a dense optical flow.

**[0041]** In order to circumvent these problems, it is proposed to use a dense optical backward flow, which is computed from the image entering the observation window  $k + 1$  to the preceding image  $k$  in the window. This approach is shown in Fig. 3(c). Rounding the optical backward flow components  $u$  and  $v$  to the nearest integer for the horizontal and vertical direction and clipping components pointing outside the valid image area, the initial mapping of pixels to cluster centers

of the new image  $k + 1$  denoted as  $\hat{\sigma}_{x,y,k+1}^{init}$  can be deduced from the previous mapping after  $L$  iterations of image  $k$  as follows:

$$\hat{\sigma}_{x,y,k+1}^{init} = \hat{\sigma}_{x+u,y+v,k}^L.$$

**[0042]** This approach solves the problem mentioned above and leads to an unambiguous label for each pixel. The only issue left, which also exists for the forward optical flow, is that the projected superpixels can be fragmented, i.e. they are not spatially coherent. In the case, in which a superpixel is fragmented, the largest fragment is determined and the other fragments are set to unassigned and are handled in the E-step of the optimization as they are part of the contour pixels.

**[0043]** In [6] a heuristic was introduced to encounter structural changes in the video volume, which are e.g. occlusions, disocclusions, and objects approaching the camera as well as zooming. The decision to split or terminate a temporal

superpixel was made based on a linear growth assumption of the superpixel size. Additionally, a separate balancing step was performed to keep the number of superpixels per frame constant. In order to simplify the treatment of the superpixels, these two steps are replaced with a single one by introducing an upper and lower bound for the superpixel size. Superpixels that are larger than the upper bound after the optimization are split and the ones that are smaller than the lower bound after the optimization are terminated. Therefore, these bounds are coupled to the number of superpixels initially specified by the user. Thus, the user defines a minimum and maximum number of superpixels per image  $N_{min}$  and  $N_{max}$ , respectively. Based on that, the upper and lower bound  $A_{low}$  and  $A_{up}$  are derived as follows

$$A_{low} = \frac{|P|}{N_{max}} \quad \text{and} \quad A_{up} = \frac{|P|}{N_{min}}$$

where  $|P|$  is the number of pixels per image. In the present implementation a number of superpixels is specified as  $N$

and  $N_{min}$  and  $N_{max}$  are set to  $\frac{1}{2}N$  and  $2N$ , respectively.

**[0044]** To summarize the proposed approach, the hybrid clustering scheme is extended to process in the *assignment* step only contour pixels in the mutable images of the sliding window. The mutable images of the sliding window are the *current* image and the *future* images.

**[0045]** New images that enter the sliding window are initialized in the following way. Between the new image that is to enter the sliding window and the latest image in the sliding window the backward flow is computed. Based on this backward flow for each pixel in the new image the corresponding pixel in the latest image is determined. Then, each pixel in the new image is assigned to the same cluster the corresponding pixel in the latest image is assigned to. The only issue left is that the projected superpixels can be fragmented, i.e. they are not spatially coherent. In the case of a fragmented superpixel the largest fragment is determined and the pixels belonging to the other fragments are set to unassigned. The unassigned pixels are handled in the E-step of the optimization.

**[0046]** Assignment changes are only done if the spatial coherency of the superpixels is guaranteed, which prevents that fragments of a temporal superpixel are split-off during the optimization. In addition, any unassigned contour pixel is assigned to the cluster of one of its adjacent pixels. As a consequence, no additional post-processing step is required to ensure the spatial coherency.

**[0047]** The contour evolution-based clustering on the video volume including the contour propagation for the initialization leads to a high temporal consistency and a high boundary/contour accuracy, even on datasets with different kinds of camera movement, non-rigid motion, and motion blur. Moreover, due to the selective processing (only contour pixels are processed during the assignment step) a competitive run-time performance is achieved.

**[0048]** A method according to the invention for generating temporally consistent superpixels for a sequence of images is schematically illustrated in Fig. 5. In a first step a cluster assignment is generated 10 for a first image of the sequence of images, e.g. by clustering pixels of the first image into superpixels or by retrieving an initial cluster assignment for the first image and processing only contour pixels with regard to their cluster assignment. Subsequent images are initialized 11 based on a label propagation using backward optical flow. Then, for subsequent images of the sequence of images, only contour pixels are processed 12 with regard to their cluster assignment.

**[0049]** Fig. 6 schematically illustrates an apparatus 20 according to the invention for generating temporally consistent superpixels for a sequence of images. The apparatus 20 comprises an input 21 for receiving a sequence of images. A cluster assignment generator 22 generates 10 a cluster assignment for a first image of the sequence of images, e.g. by clustering pixels of the first image into superpixels or by retrieving an initial cluster assignment for the first image and processing only contour pixels with regard to their cluster assignment. A label propagator 23 initializes 11 subsequent images based on a label propagation using backward optical flow. A contour pixel processor 24 then processes 12 only contour pixels with regard to their cluster assignment for subsequent images of the sequence of images. The resulting superpixels are preferably made available at an output 25. Of course, the different units 22, 23, 24 of the apparatus 20 may likewise be fully or partially combined into a single unit or implemented as software running on a processor. Also, the input 21 and the output 25 can likewise form a single bi-directional interface.

## REFERENCES

**[0050]**

[1] X. Ren et al.: "Learning a classification model for segmentation", 2003 IEEE International Conference on Computer Vision (ICCV), pp. 10-17.

[2] O. Veksler et al.: "Superpixels and Supervoxels in an Energy Optimization Framework", Proceedings of the 11th European conference on Computer vision: Part V (ECCV'10) (2010), pp. 211-224.

[3] A. Levinstein et al.: "Spatiotemporal Closure", Proceedings of the 10th Asian conference on Computer vision: Part I (ACCV'10) (2010), pp. 369-382.

[4] A. Schick et al.: "Measuring and Evaluating the Compactness of Superpixels", 21st International Conference on Pattern Recognition (ICPR) (2012), pp. 930-934.

[5] R. Achanta et al.: "SLIC superpixels compared to state-of-the-art superpixel methods", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 34 (2012), pp. 2274-2282.

[6] M. Reso et al.: "Temporally Consistent Superpixels", 2013 IEEE International Conference on Computer Vision (ICCV), pp. 385-392.

## Claims

1. A method for generating temporally consistent superpixels for a sequence of images, the method **comprising**:
  - generating (10) a cluster assignment for a first image of the sequence of images;
  - initializing (11) subsequent images based on a label propagation using backward optical flow; and
  - for subsequent images of the sequence of images, processing (12) only contour pixels with regard to their cluster assignment.
2. The method according to claim 1, **wherein** generating (10) the cluster assignment for the first image of the sequence of images comprises clustering pixels of the first image into superpixels.
3. The method according to claim 1, **wherein** generating (10) the cluster assignment for the first image of the sequence of images comprises retrieving an initial cluster assignment for the first image and processing only contour pixels with regard to their cluster assignment.
4. The method according to one of the preceding claims, **wherein** an assignment change of a pixel is only done if the spatial coherency of the superpixels is guaranteed.
5. The method according to one of the preceding claims, **wherein** in case of fragmentation of a superpixel, the largest fragment of the superpixel is determined and pixels belonging to the remaining fragments of the superpixel are set to unassigned.
6. The method according to claim 5, **wherein** any unassigned pixel is assigned to the cluster of one of its adjacent pixels.
7. The method according to one of the preceding claims, **further** comprising splitting superpixels that are larger than an upper bound and terminating superpixels that are smaller than a lower bound.
8. The method according to claim 7, **wherein** the upper bound and the lower bound are determined from the number of pixels per image and a minimum and a maximum number of superpixels per image, respectively.
9. An apparatus (20) configured to generate temporally consistent superpixels for a sequence of images, **wherein** the apparatus (20) comprises:
  - a cluster assignment generator (22) configured to generate (10) a cluster assignment for a first image of the sequence of images;
  - a label propagator (23) configured to initialize (11) subsequent images based on a label propagation using backward optical flow; and
  - a contour pixel processor (24) configured to process (12) only contour pixels with regard to their cluster assignment for subsequent images of the sequence of images.
10. A computer readable storage medium having stored therein instructions enabling generating temporally consistent



superpixels for a sequence of images, wherein the instructions, when executed by a computer, cause the computer to:

- generate (10) a cluster assignment for a first image of the sequence of images;
- initialize (11) subsequent images based on a label propagation using backward optical flow; and
- for subsequent images of the sequence of images, process (12) only contour pixels with regard to their cluster assignment.

11. A computer program comprising instructions enabling generating temporally consistent superpixels for a sequence of images, wherein the instructions, when executed by a computer, cause the computer to:

- generate (10) a cluster assignment for a first image of the sequence of images;
- initialize (11) subsequent images based on a label propagation using backward optical flow; and
- for subsequent images of the sequence of images, process (12) only contour pixels with regard to their cluster assignment.







Fig. 1




**Fig. 2**




**Fig. 3**



**Fig. 4**



**Fig. 5**



**Fig. 6**



## EUROPEAN SEARCH REPORT

Application Number  
EP 15 17 2032

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
A	<p>Rudolf Mester ET AL: "Multichannel Segmentation Using Contour Relaxation: Fast Super-Pixels and Temporal Propagation"</p> <p>In: "LECTURE NOTES IN COMPUTER SCIENCE", 1 January 2011 (2011-01-01), Springer Berlin Heidelberg, Berlin, Heidelberg, XP055124865, ISSN: 0302-9743 ISBN: 978-3-54-045234-8 vol. 6688, pages 250-261, DOI: 10.1007/978-3-642-21227-7_24, * abstract *</p> <p>* sections 4, 6 *</p> <p>-----</p>	1-11	INV. G06T7/00
A,D	<p>RESO MATTHIAS ET AL: "Temporally Consistent Superpixels", 2013 IEEE INTERNATIONAL CONFERENCE ON COMPUTER VISION, IEEE, 1 December 2013 (2013-12-01), pages 385-392, XP032572909, ISSN: 1550-5499, DOI: 10.1109/ICCV.2013.55 [retrieved on 2014-02-28] * abstract *</p> <p>* sections 4.3, 4.4, 4.5 *</p> <p>-----</p>	1-11	<p>TECHNICAL FIELDS SEARCHED (IPC)</p> <p>G06T</p>
A	<p>IVANA MIKIC * ET AL: "Segmentation and Tracking in Echocardiographic Sequences: Active Contours Guided by Optical Flow Estimates", IEEE TRANSACTIONS ON MEDICAL IMAGING, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 17, no. 2, 1 April 1998 (1998-04-01), XP011035719, ISSN: 0278-0062 * abstract *</p> <p>* section II *</p> <p>-----</p> <p>-/--</p>	1-11	
The present search report has been drawn up for all claims			
Place of search		Date of completion of the search	Examiner
Munich		12 October 2015	Eveno, Nicolas
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons</p> <p>&amp; : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.82 (P04C01)



## EUROPEAN SEARCH REPORT

Application Number  
EP 15 17 2032

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
A	KRIM H ET AL: "Fast Incorporation of Optical Flow Into Active Polygons", IEEE TRANSACTIONS ON IMAGE PROCESSING, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 14, no. 6, 1 June 2005 (2005-06-01), pages 745-759, XP011131845, ISSN: 1057-7149, DOI: 10.1109/TIP.2005.847286 * abstract * * section I *	1-11	
			TECHNICAL FIELDS SEARCHED (IPC)
The present search report has been drawn up for all claims			
Place of search		Date of completion of the search	Examiner
Munich		12 October 2015	Eveno, Nicolas
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 03.82 (P04C01)

## REFERENCES CITED IN THE DESCRIPTION

*This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.*

## Non-patent literature cited in the description

- **X. REN et al.** Learning a classification model for segmentation. *IEEE International Conference on Computer Vision (ICCV)*, 2003, 10-17 [0050]
- **O. VEKSLER et al.** Superpixels and Supervoxels in an Energy Optimization Framework. *Proceedings of the 11th European conference on Computer vision: Part V (ECCV'10)*, 2010, 211-224 [0050]
- **A. LEVINSHTEIN et al.** Spatiotemporal Closure. *Proceedings of the 10th Asian conference on Computer vision: Part I (ACCV'10)*, 2010, 369-382 [0050]
- **A. SCHICK et al.** Measuring and Evaluating the Compactness of Superpixels. *21st International Conference on Pattern Recognition (ICPR)*, 2012, 930-934 [0050]
- **R. ACHANTA et al.** SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2012, vol. 34, 2274-2282 [0050]
- **M. RESO et al.** Temporally Consistent Superpixels. *IEEE International Conference on Computer Vision (ICCV)*, 2013, 385-392 [0050]